Low-Cost, Light-Weight Mars Landing System

Richard W. Warwick
Lockheed Martin Space Systems Company
P.O. Box 179
Denver, CO 80201
richard.w.warwick@lmco.com

Abstract— Based on the Mars 2001 lander, a design for a low cost Mars landing system with a 50-150 kg payload class capability is presented. Discussion includes a history of the design considerations to achieve high payload masses, including technology insertion, as well as discussion of design optimization. The enhancements and changes based on lessons learned from MPL are presented, as well as additional options for addressing issues raised during the failure review investigations. Also addressed is the adaptability of the design to different payload configurations and launch opportunities. Operational approaches to accommodate large number of payloads, onboard resource allocations and budgeting issues, timeline and planning constraints imposed by round trip light time, communication coverage, and shift planning are discussed. Parameters unique to landed Mars missions that can be traded against one another to achieve specific mission objectives are also discussed along with observations on their impact to mission design and planning.

TABLE OF CONTENTS

1. INTRODUCTION

Based on the development work of the Mars Surveyor Program's Mars Polar Lander and 2001 lander, a low cost, light-weight Mars landing system has been developed with substantial payload capability. By leveraging the development of the Mars Surveyor lander series, a Mars lander flight system weighing no more than 700 kg at Earth launch can deliver a payload of over 100 kg to the Martian surface. This paper discusses some of the history, design trades and operations considerations of such a landing system and a large payload suite.

The Mars Polar Lander was designed to carry 24 kg of payload to Mars and weigh barely 600 kg at Earth launch. The 2001 lander was upgraded to carry 66 kg of payload to Mars for 687 kg of Earth launch mass. The design can be upgraded to carry up to 120 kg of payload, depending on the arrival conditions at Mars, for approximately 700 kg of Earth launch mass. This low mass at Earth launch can enable the use of the low cost Med-Lite launch vehicle class, depending on the Mars launch opportunity.

2. MSP 2001 PROJECT HISTORY

The Mars Surveyor Program 2001 project was the second procurement in the Mars Surveyor Program series. Mars Global Surveyor was started prior to the inception of the MSP (although operationally it became part of the Mars Surveyor Operations Project, the MSOP). The MSP planned to send an orbiter and a lander to Mars every launch opportunity through at least 2005. The MSP'98 project, which ultimately became Mars Climate Orbiter and Mars Polar Lander, was the first procurement. As originally proposed, the 2001 orbiter and lander were technological and performance upgrades to the MSP'98 spacecraft. Planning for the 2001 orbiter began in early in 1996, while preparations for the lander did not begin until late in 1996, spurred primarily by the announcement that evidence for fossil life had been found in ALH80001.

Design Requirements

Although the original procurement for the MSP spacecraft addressed the capabilities for multiple launch opportunities, the MSP'98 project quickly became focused on just the 1998 Mars launch opportunity. Several things made it quickly apparent that the 2001 spacecraft would not be a build-to-print copy.

[1] 0-7803-7651-X/03/$17.00 © 2003 IEEE
[2] IEEEAC paper #1473, Updated January 28, 2003

The launch opportunity itself meant both launch and arrival conditions would be different. This meant differences in TPS on the heatshield, differences in available power, etc. Also, because the landing site desired for the payload was equatorial instead of polar, the thermal environments would be more extreme, in addition to significantly longer nights. The payload mass was also substantially greater for the ’01 lander than for MPL, nearly a factor of 3 (24 kg versus 66 kg). Finally, the simple fact that the projects were treated more as discrete projects than as a continuous program meant that some requirements were different from MSP’98 simply because there was no enforcement of commonality.

Another demanding requirement on the ’01 lander flight system was operations time. Several of the instruments required continuous, 24.6 hour per sol operations. A major trade study was conducted early in 1997 to determine the feasibility of allowing payloads to operate all day and all night. While MPL had the option of powering up just about whenever it wanted to because of the continuous availability of solar power during the south pole summer, the long, 12-hour nights at the equator would have demanded each battery be at least 40 Amp hours and the arrays be 50% larger than the 4.2 m^2 they already were. The spacecraft was nominally designed for 7 hours of “awake” operations, with the C&DH fully powered, and 17 hours of “sleep mode” where the lander computer was unavailable. The alternative was that if a payload could buffer its own data, and make any sequencing changes without the lander computer, the only constraint would then be total energy usage and total daily data volume generation (because of the total downlink budget). The next size batteries that most vendors had experience with were 50 A-hr batteries, and they were simply too large to accommodate. The flexible arrays were also as large as could be packaged. As a result, changes were made to accommodate a larger number of switches that would remain on when the lander powered off and payloads that could operate without the lander computer would have no time limitations on their operation. The operations sequence also had to be modified to have the lander wake up for a few minutes every 2 hours to check the battery state of charge to be sure that no payload had shorted at the fuse limit and was draining the battery. It also meant that data could be dumped from the instruments every two hours so they could reduce the amount they needed to buffer themselves. This feature, the ability to operate the payloads without the presence of the lander computer, turned out to be such a benefit, that its use has been proposed for other missions with critical events (such as flyby’s) to ensure that data can be returned even if the spacecraft were to take an SEU or have a processor reset for some other reason.

Design Optimization

The MSP’98 lander design requirements placed greater restrictions on the spacecraft than on Mars Pathfinder. MPL and Pathfinder (counting Sojourner) carried about the same payload mass, but MPL was given a requirement early on to fit on a Delta 7325, which had about half the lift capability of the Delta 7925 that launched Pathfinder. The later addition of the DS-2 microprobes bot an extra SRM, making the launch vehicle a 7425, but the total launch capability was still significantly less than Pathfinder actually weighed at lift-off. This meant that the landing system had to have a greater payload fraction than Pathfinder.

The MSP’98 program not only achieved this greater payload fraction, but the lander actually came in significantly lighter than expected. The combination of this, plus the technology insertion used for the 2001 lander, permitted the 2001 lander payload mass to be increased to 66 kg. In fact, at the time the program was placed on hold, the ’01 lander was projected to come in nearly 18 kg light.

There were several ways that the 2001 program increased the payload mass. While some was achieved by technology insertion, some was achieved by further optimizing the design. For example, MSP’98 had four pairs of bipod struts that held the lander to the backshell. MSP’01 used three instead. The attachments between these bipods and the decks eliminated the guide rails used by MSP’98, not only eliminating a potential failure mode, but also resulting in a lighter structure (see figure 1.).

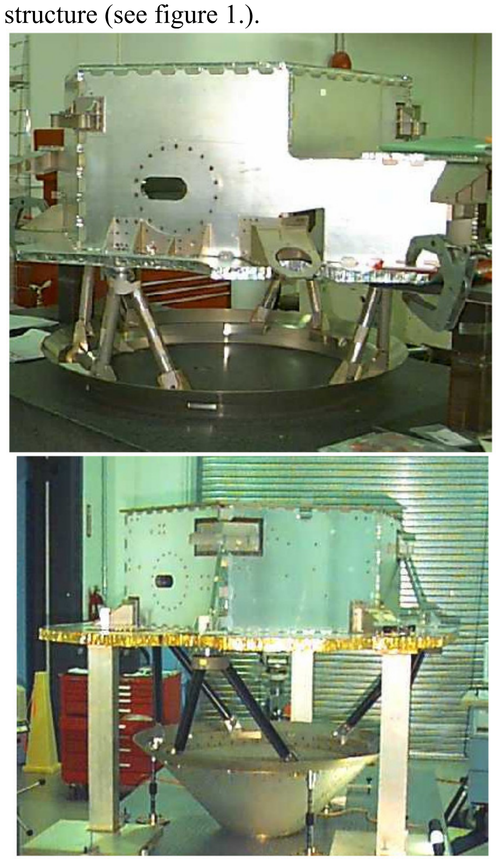

Figure 1. MPL (top) structure and MSP'01 (left) structure showing optimization by increasing available volume while decreasing mass of bipods.

The 2001 lander also benefited from the MSP'98 program by realizing actuals early enough in the program. For example, after the MPL descent thrusters had been tested, the '01 project took advantage of this more advanced state of knowledge of performance. Many other components that were weighed from the MSP'98 program that would fly on '01 could be reported with lower individual margins on them. Because the lander was the second generation spacecraft, it benefited from the development efforts of both MSP'98 and MSP'01.

The last decision made to reduce the landed mass of the flight system and make more mass available to payloads was the decision to eliminate a direct-to-earth telecom link. At the time the decision was made, MGS was in orbit, and both MCO and Odyssey were expected to be on station during the lander mission. With three potential relay links, it was deemed unnecessary to fly a landed X-band system. This freed up another fifteen or more kilograms for payload use.

Technology Insertion

Despite all the benefits of being the second generation spacecraft, more performance was still needed, especially early in the program before some of the benefits from being second were realized. Other considerations also drove the technology programs, such as the longer nights, wider thermal ranges and greater payload mass and power needs.

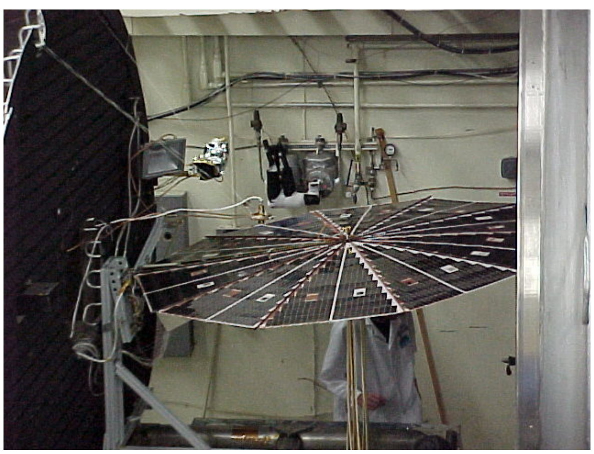

Figure 1. 2001 lander -Y flight solar array entering thermal-vacuum testing.

Several new technologies were chosen that addressed multiple new requirements for the '01 lander. In the end, only one was implemented because of mass alone. Rather than using the MPL and Pathfinder parachute design, the '01 project opted to change the disk-gap-band ratio back to the Viking values. Pathfinder had changed the parachute because the airbag lander could not tolerate the high amplitude swings on the Viking parachute design and MPL had simply used the Pathfinder parachute as it was. The '01 lander needed to land more mass, however, and the Viking disk-gap-bad ratio parachute provided more drag for the same parachute area, so the change was made.

The electrical power subsystem probably incurred the most change. First, because the '01 lander needed to collect more energy over a diurnal cycle, the array area had to be increased. This would not only increase the mass of the arrays, but make them larger, and packaging was already a problem, even though '01 was already using a Pathfinder sized 2.65m aeroshell instead of the 2.4 meter aeroshell used by MPL. A trade was conducted early in the program, and a vendor selected to provide an array that packaged in a much smaller volume than the '98 arrays, and weighed less. The result was an array that produced (at one-sun, AM0) nearly 120 W/kg. The lightweight flexible arrays were delivered, but never installed on the lander (see figure 2.). Part of the trade in 1997 for the flexible solar arrays also considered the cells to use. LILT testing showed that high efficiency silicon cells (flown on Stardust) performed slightly better than regular GaAs. Because the process for laying down silicon cells on the flexible gores had been qualified, but GaAs had not at the time, the decision to go with hi efficiency silicon was an easy one, and they were incorporated into the design.

Also, because of the longer nights, more battery capacity was needed. The temperature limits of the Nickel Hydrogen battery flown on MSP'98 were among the tightest of any component, and so a Li-Ion battery was selected for the '01 lander. The single drawback of the Li-Ion battery was that an open cell was a credible failure, thus two batteries would be needed. However because the energy density by mass and volume were three and two times that of the Nickel Hydrogen battery, respectively, the '01 lander had three times as many amp hours as '98 for slightly less mass and volume (see figure 3.). The single fault tolerance policy meant that, while for certain battery related faults the '01 lander could count on 25 A-hrs as opposed to MPL's 16 A-hrs, other fault scenarios allowed the '01 lander to count on 50 A-hrs. The performance of the Li-Ion batteries was also such that only at –20 deg C was the capacity limited to 25 A-hrs. While design and analysis never took advantage of more than 25 A-hrs, performance at the temperatures the battery was likely to operate at provided 35 A-hrs, providing additional margin under expected landed ops conditions.

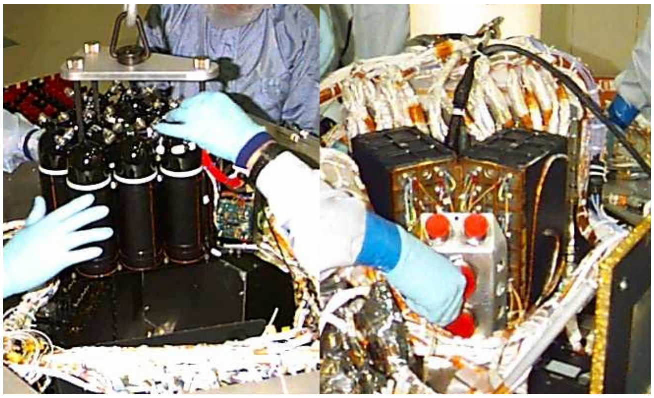

Figure 2. MPL's single (left) 16-Ahr nickel-hydrogen battery and MSP'01 lander (right) dual 25 A-hr lithium-ion batteries take up approximately the same volume.

Another change to the power system was use of HDI modules on the power switch cards. While the original design was not the one ultimately implemented, the end result was still more than double the switches for the same number of cards. New high efficiency power supplies were developed to reduce overall power consumption. The new power supplies also eliminated a card in the C&DH. Mars Odyssey is presently flying both these new technologies.

Additional Enhancements

Since the cancellation of the '01 lander launch, additional studies have been conducted to further optimize the capacity of the design, and one in particular bears mentioning. A trade conducted during phase B in 1997 was to replace the pulse mode terminal descent propulsion system with a throttled propulsion system. The idea envisioned using MR-104 thrusters instead of MR-107 thrusters. Three sets of three would each be manifolded together to provide a set-up similar to Viking, but without the expensive special thruster. It was ultimately abandoned in 1997 due to cost and schedule for an '01 launch. However, in 2002, an analysis was done of a similar system, but using the original Viking engine instead of the MR-104's proposed in 1997. Monte Carlo analysis showed that the resultant payload mass of the lander would be conservatively 120 kg; nearly double the 66 kg design requirement. The reason for the increase is that the three Viking engines had significantly more thrust than the 12 MR-107's, so the burn would have been more impulsive. Velocities coming off the parachute are typically 80-100 m/s, but the delta-V produced by the MR-107's was well over 200 m/s, indicating most of the burn was being consumed by gravity losses. Total burn times in the Viking engine Monte Carlo were about half that of the MR-107 system, showing that indeed the gravity losses were being dramatically reduced. Further optimization could be realized by re-examining the constant velocity phase of terminal descent during the last 10-12 meters.

Since the trade studies in 1997 on solar array design, considerable advances as well as LILT testing under Mars filters on triple junction GaAs cells has shown that they would provide even more power on the same array, although the mass would increase. MER is making use of these cells and will fly them to Mars in 2003.

3. MPL Failure and Implications

Ultimately, the loss of MPL, without telemetry to provide positive indication of the failure, resulted in the cancellation of the 2001 lander. While the most probable cause was identified, there were seven failure modes that were considered plausible. Because the other six are at least an order of magnitude or even several orders of magnitude less likely to have occurred, they are not discussed here. See Reference 1 for more details on all identified failure modes. Reference 2 provides greater detail on the most probable cause; it is only summarized here.

Most Probable Cause

On January 18, 2000, during an EDL run on the '01 STL, the test engineer accidentally entered the command to simulate the touchdown signal from the touchdown sensors early. He quickly entered a new command to remove the simulated presence of the signal, but the simulation still resulted in a crash landing. Investigation of the problem revealed that even though the simulated touchdown sensor signal had been removed, the touchdown flag had still been set and never cleared, and the engines were cut off at approximately 40 meters above the ground.

A subsequent investigation (see Reference 3) later revealed that, while there were many contributing factors, three independent things conspired to produce this software problem. It is interesting to note that it took all three of these situations to produce the fatal flaw; had any one of them not been present, the lander would have made it past this point in EDL. The first was a faulty requirement flowdown. During a trade study in 1996, Hall effect sensors, used successfully on many missions, were chosen to provide indications of touchdown. The mechanisms engineers knew that transients could result from these sensors (one was seen during a leg deployment test in 1997, but this was expected), and communicated that a requirement not to use the signals until the lander was close to the ground was needed. The resulting requirement was a "shall not". Because of the inherent difficulty in translating a shall not, the intent of the shall not was not captured when the requirements were flowed down to the FSW requirements specification. Also, the requirement to not use the signals until the lander was close to the ground (originally 12 meters, and later revised to 40, which was where the wide beam radar lost velocity lock), was not identified as mission critical and thus not tracked in the DPD and no resulting system level verification requirement was generated. This is why, when the leg was found to be wired wrong during an EDL test on the spacecraft, the test was not re-run. All objective of that test had been met.

However, this requirements translation alone was not enough to cause the fatal flaw. It took the combination of all three things to produce the failure mode. The second requirement that contributed to the loss of MPL was a requirement that all FSW objects be started prior to entry. The reason for this requirement was to reduce the risk of a flight processor reset during EDL (which would be fatal) caused by a new object starting and causing a momentary spike in CPU utilization. In hindsight, it is believed that this small object would not have caused a problem, but the requirement applied to all FSW objects. The result was that the object would now be running when the legs were deployed, although FSW would not examine the flag prior to reaching 40 m above the ground.

Even this was not enough to cause the fatal flaw. A new requirement, added in 1998 to put the touchdown flag into channelized telemetry, required a code change to the touchdown sensor code. The earlier version of the code would have been insensitive to the transients. The new version, however, did not clear the flag (there was no requirement to do so), and the stage was now set for the failure mode that most likely resulted in the loss of MPL.

The study of this failure mode shows the difficulties in uncovering subtle situations that can be mission critical. If any one of the three conditions had not existed, this failure mode would not have occurred. Had the "shall not" been written as a proper shall (e.g. the flight software shall reject transients from the touchdown sensors), there would have been a test case to verify that transients were rejected. Had the flight software object not been started until radar cutoff at 40 meters, it would not have set the flag from the leg deploy transient signal. Had the change to put the touchdown flag into telemetry not been made, the code would not have been changed, and the older version that was impervious to the error would have been used.

Testing both before and subsequent to uncovering this flaw resulted in a range of statistics on the probability of this error occurring. The transients appeared to be statistical, and no systemic behavior was observed (one leg didn't always produce extra long transients). The touchdown sensor code read the inputs from the leg sensors every 10 ms, and two consecutive positive reads would indicate a touchdown by design. Several tests produced no transient signals at all. These are potentially suspect because the sensors are expected to always produce at least a short transient due to the mechanical motion of the deploying leg locking into position, so counting the zero signal duration transients is arguable. Some tests were conducted at an angle to simulate the load conditions that would be experienced in Mars gravity during EDL, and some were conducted in thermal-vacuum facilities with Mars atmospheric pressures. Some were conducted with the lander hard-mounted to a fixture and some were conducted with the lander suspended from a cable. Out of 47 leg deployment tests, with several being suspect, the statistics range from a 47% chance per leg to a 93% chance per leg of producing a transient long enough to set the touchdown flag (see figure 4.). After 3 leg deploys, this translates to between 86.7% to 99.96% chance that MPL experienced this failure mode. MPL had, at best, only a 1 in 8 chance of getting past this error. While there was no telemetry during EDL to verify this failure mode actually occurred, its likelihood is orders of magnitude greater than any other identified failure mode.

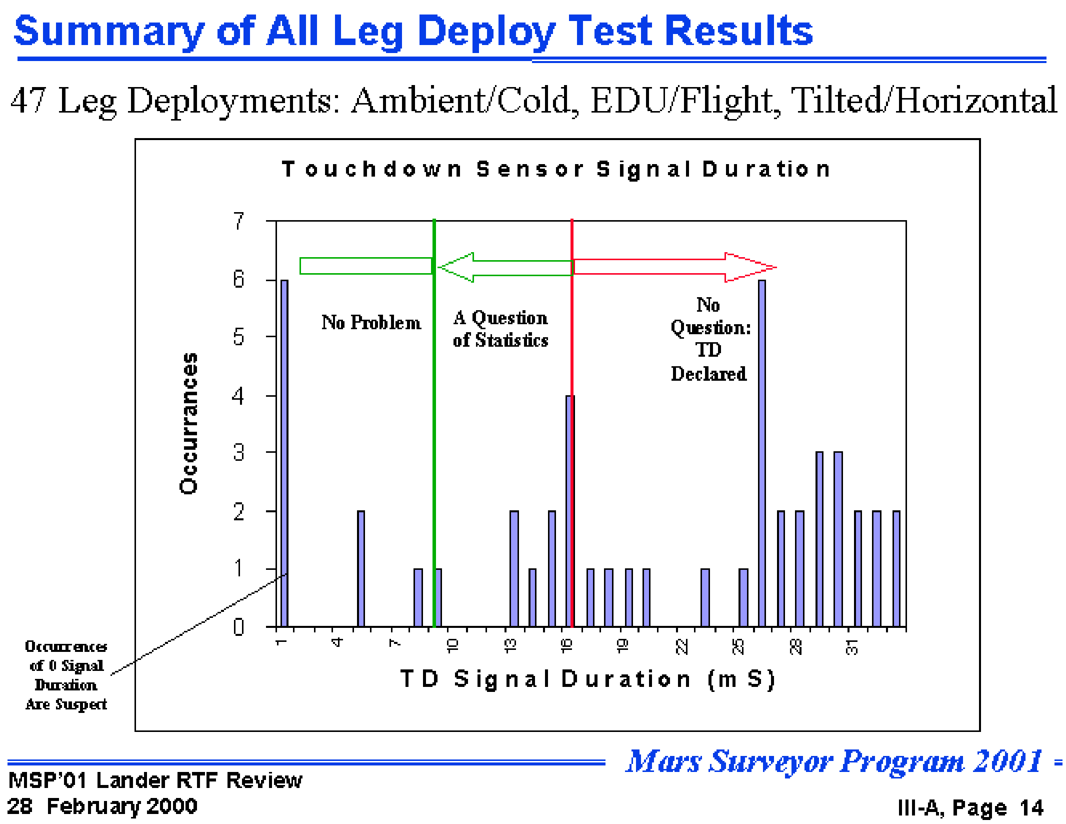

Figure 3. Transient signal durations from touchdown sensors of all leg deployment tests.

The result would have been premature thrust termination at an altitude of 40 m. The vertical speed would have been about 13 m/s, and in the 2.3 seconds it would take MPL to reach the surface, the speed would increase to almost 22 m/s. The resultant impact would have been over 70,000 J, rather than the 1000-2000 J the lander would have seen had it touched down at the programmed velocity. It is interesting to note that because the time to fall was so short, the lander would have impacted in almost the exact orientation it was supposed to land. Assuming the lander was properly controlling the rates about each axis to the specified levels (and in most cases, they would be much lower), it would not have been able to rotate more than a few degrees about any axis during the fall. It is also interesting to note that, depending on the surface properties and the manner in which the structure deformed and crumpled, many of the avionics components may have survived the crash because of the g-loads they were qualified to for launch and pyro shock environments. One piece of hardware that was almost certainly a casualty, however, would have been the internal sidewall to which most of the telecom waveguide hardware was mounted. Even if the computer had survived, there would have been no way to transmit.

Implications for the 2001 Lander

The touchdown sensor code error that most likely caused the loss of MPL was fixed and tested on MSP'01 just a few days after its discovery, and in fact the corrected version was delivered with the second flight software build, which represented nearly half of the 400 hours of test on the spacecraft prior to the lander being placed on hold. However, the lack of positive indication of the MPL failure mode has meant that all the potential MPL failure modes be examined and potentially addressed if they are applicable to the '01 lander. There are a few that are not applicable (two of the MPL failure modes were because the A side transponder was a Cassini spare that had a flag against it, obviously not applicable to the '01 lander's new transponders). In addition, several review boards have recommended changes that, while not directly addressing a failure mode, address other perceived shortcomings in the Mars Surveyor Program.

In this last category, two primary recommendations were to add EDL telemetry and add a DTE X-band link. Both MPL and '01 projects chose not to implement EDL telemetry. Such a system would actually decrease the chances of a successful landing by using valuable resources (power and CPU margin) during the critical EDL phase. That coupled with the cost savings to each project made it an easy decision for each of the projects. Had the MSP been a single, more coherent program, rather than a collection of individual projects, especially in light of the fact that each spacecraft was to be an upgrade of its predecessor, it is likely that a requirement for critical event communications would have been levied.

The X-band link that was deleted from the '01 mission in 1997 can be easily re-incorporated. In fact, the resources to add it back in had always been held, with the exception of the mass that was freed up by its elimination, and that mass has been regained because the lander is lighter than originally projected.

Other recommendations are tied to other identified failure modes. They are not discussed in detail here. The return to flight review identified mitigation strategies for every issue raised by the MPL failure (see Reference 4). Ultimately, it was decided that too little time existed (less than 15 months) to implement some of them, but by launching the lander in a later window, the schedule constraint can be alleviated, and the lander can be safely flown with upgrades to address all pertinent, identified MPL failure modes.

4. ACCOMMODATION OF LARGE PAYLOAD SUITES

The primary challenge faced by the '01 project was a payload suite that was always, in every way, demanding more resources than the payloads had on MPL: more mass, more volume, more power, more switches, more analog and digital I/O, more memory. Every spacecraft resource had to be increased or more of the existing ones were allocated to payloads. But it was the mass and volume demands that placed the greatest challenges on the design.

Ultimately, MPL and the '01 lander had very little functional duplication across the flight elements, aside from the A-side/B-side redundancy built in by design. A single propulsion system was used both for cruise and EDL. This is one reason the launch stack of both MPL and '01 was lighter than Pathfinder, which used a propulsion system during cruise in space, and a completely separate airbag and rad rocket system for terminal descent. There were no separate electronics boxes discarded with the aeroshell. All cruise stage hardware interfaced directly with the lander C&DH inside the aeroshell. The cruise stage was really just a launch adaptor ring with solar arrays, star cameras and antennas. There were no "smarts" on it.

A lesson learned that applies in general across designs, not just those for large payloads, but that was important to realize in order to accommodate the large instrument suite in such a constrained program, was the realization that the only way to truly eliminate costs was to eliminate requirements. If a function had to be performed, something had to be thought of, designed, built, tested and all that paid for to do it. Second order savings could be realized by being clever in any of the steps just mentioned, but it was by challenging and eliminating some requirements that resources (dollars, mass, switches, etc.) were freed up to be used to accommodate the large payload suite.

Configuration

The original intent of the ’01 lander was to carry the Athena rovers and their instrument suite to Mars. These are the same payloads that have ultimately evolved to become the MER’s. At the time, the rover and its payload and lander mounted support equipment were expected to be only 45 kg. Shortly after the go-ahead for an ’01 lander, NASA decided to add payloads from HEDS to begin preparing for human missions. The large volume needed by this large payload suite drove substantial changes in the lander design from MPL. First, the size of the lander and its aeroshell were increased, to 2.65 m diameter, to provide more backshell volume and lander deck space. The entire lander primary enclosure was redesigned to increase the packaging efficiency.

There were other considerations as well. Payload fields of view and rover deployment requirements also had impacts on configuration. The UHF antenna fields also impacted payload layout, and at one point, the payload configuration demanded that the UHF be deployed upwards on a pedestal to reduce the field strengths seen at the payloads. Ranges of motion for items like arms or sample handing hardware and surface access play into the equation.

Subsequent to the lander's removal from the ’01 launch opportunity, plans to fly it in 2003 and as a scout in 2007 have resulted in evaluating many other configurations. The lessons learned can be summarized as follows:

- Deck space is easier to come by than airspace above it because of the legs, the backshell attachment structure and ultimately, the backshell. Until a new aeroshell design is developed for Mars, these types of limitations will challenge payload packagers on soft lander designs.
- De-centralized or modular payloads have a better chance of being successfully integrated than larger contiguous packages. The incursion of the parachute canister into the available volume has a dramatic impact on payload configuration. Again, this will continue to be the case until either a new aeroshell design is qualified for Mars or a new parachute deployment scheme that dramatically reduces the size of the canister becomes available.
- Constraints on CG location greatly impact configuration. The less sensitive the design can be to CG location, the greater flexibility in packaging. Landers have unique requirements imposed by descent that offer low cost opportunities to alleviate the CG problem. For example, over-sizing the descent propulsion thrusters slightly allows greater flexibility in CG while potentially increasing control margins for a relatively low cost.

Resource Allocations

Because of the large number of payloads on the ’01 lander, an early design change was made to add a second payload interface card into the C&DH. The total payload resource usage was nearly double that of MPL. Should the lander fly in the future, the likely addition of a requirement to provide positive indication of separation events will consume more of these extra resources. In 1997, payloads were discouraged from using the 1553 interface (which was originally designed into MPL to interface with the landing radar). In hindsight, if the other digital and analog I/O could have been reduced, it might have been preferable to have the payloads use the 1553 interface, given the large number of them.

Power switches and available energy are also precious resources on a landed mission. One method of making more of both available is to have instruments share resources with flight system elements that aren’t used after landing. For example, the star trackers were discarded with the cruise stage, and the radar isn’t used after EDL. The resources they used become available after touchdown, and reassigning those to payloads after touchdown is a more efficient use of valuable resources.

Attempts to accommodate large payload suites have resulted in several lessons learned unique to accommodating large payload suites, which include:

- Using bus type architectures such as 1553 rather than individual interfaces for each instrument reduces constraints on number of instruments. While there are penalties for such architectures (power draw for example), they may be reduced or even turned into benefits as the number of instruments climbs.
- Careful assessment of operational needs of each instrument can result in a larger number of instruments using the same limited resources. By grouping payloads onto common resources and operating them in such a manner that only one is utilizing that resource at a time, a larger number of instruments can be flown. This applies to switches and channels as well as power and downlink data volumes.
- For landing systems that don’t use or make minimal use of systems subject to the exponential nature of the rocket equation, the practice of leaving unnecessary hardware behind on a previous stage to save mass or some other resource can only be proven beneficial by a detailed trade study. (The MPL and ’01 EDL systems were equivalent to propulsion systems with specific impulses on the order of 1200-1500 seconds). The result is usually that there will be savings in some resources and costs in others, and they must all be identified and weighed together from a systems viewpoint to make the best decision. While MPL

and the '01 lander ultimately flew most hardware to the ground and discarded very little prior to and during EDL, subtle changes in requirements could have easily changed the outcome.

5. Constraints on Operations of Landed Payloads

Orbital Mechanics

The Martian solar day (sol) is about 24.6 hours. Because of this, time on Mars slowly walks out of phase with time on Earth. This is a problem when planning shifts. Take a simple example, where a single 8-hour shift per day is planned to prepare commands and analyze data (this is an unrealistically small amount of coverage, but it serves to illustrate the point). To use the same people, they have to keep coming to work 0.6 hours later each day. Every two weeks or so, they would have moved completely to the next shift.

Telecommunications

There are several constraints placed on landed Mars mission because of telecommunications. First, a mission may have to compete with other missions for DSN time. One can assume that a 90 sol landed Mars mission would get first priority (but what if there are two or more Mars missions on the surface at once). Such scheduling concerns are very real right now with the armada of spacecraft headed for Mars in the next year.

The first problem with telecommunications is what kind of link is available. Does the lander have only a Direct-To-Earth link (DTE), does it have only a relay system for use with an orbital asset, or does it have both. In general, an X-band DTE system will have, at least by design, anywhere from a factor of four to an order of magnitude less performance in just about every measurable way. With such performance, it is arguable they should not be included in the design. The significant advantage is that the communication coverage is continuous as long as Earth is visible in the sky. Also, while you can't count on it during the design phase, the actual performance of the system will usually be a factor of two to four better than planned, so in practice, it won't be quite as limiting. The big limitations are total data volume. This is constrained on one end by lower data rates than relay links, usually by several orders of magnitude. For example, UHF relay rates on Odyssey are 128 kbps and 256 kbps (MRO will be capable of even higher rates), while planned rates for X-band DTE links from the surface are generally 1600 to 2400 bps. Actual performance (such as from Pathfinder) was closer to 8000 bps. So, to get the same amount of data (in the design phase) from a DTE at 2400 bps versus a UHF relay at 128 kbps, it would take, assuming the UHF pass were five to ten minutes (the time an orbiter at 400 km altitude would be visible), the DTE would have to operate nearly nine hours. If the system actually performed at 9600 bps upon arrival, it would still take over two hours to get the data back. In general, the additional power dissipated during the transmit period would put increased demands on the thermal subsystem to dump the extra heat. Given that the general direction of the thermal design is to hold in the heat during the nighttime, this usually means either active heat rejection systems on the surface, or severe limitations on DTE transmit times. A UHF link could dump the same amount of data to an orbiter in ten minutes or less, and at a typical Odyssey X-band rate of nearly 40 kbps, that data would be back to Earth in half an hour. Assuming full cost accounting, (in other words the orbiter charges the lander for that DSN time), the relay link can save substantial money over the DTE. If the lander is not charged for the half hour of orbiter downlink time, the cost savings to the lander project itself are even greater.

The drawback of relay operations is that except for landing sites near the pole, a lander mission can expect, on average, two passes per sol for each orbital asset (assuming the orbiter is in a polar orbit, as most are and plan to be). Polar lander missions can expect as many passes as twelve per sol per orbiter. Only having two passes per sol can add serious risk to landed operations. See figure 5 for an example.

Assume an orbiter is in a 4 PM orbit, the lander is near equatorial, so the lander gets a UHF pass every sol at about 4 PM and about 4AM. It may not be a full pass, as the combination of the orbiters true anomaly and the rotation of the landing site can conspire to limit the time the orbiter is visible to something less than the maximum possible pass, which would be an orbiter passing directly overhead, but assume two links (there are sols where only one is available because of this geometry problem, and others where 3 passes occur). Say an anomaly happens at noon Mars time today. At 4 PM, a relay pass sends that data to an orbiter, and within an hour, the telemetry has been displayed to the operators on Earth. At this point, the pass between the orbiter and the lander has ended, so nothing can be sent to the lander until the next morning at the earliest. Even if a scheme were envisioned such that anomalous data were relayed immediately upon receipt by an orbiter, light time alone between Earth and Mars would prevent any action from being taken. So the data is analyzed, a solution determined, and new commands ready to go. Assuming there is either continuous coverage of the orbiter or at least a chance to radiate the commands prior to the next lander pass at 4 AM, the new commands would be sent to the orbiter and relayed to the lander at 4 AM. (If the orbiter coverage is such that there is only, say, one 8-hour pass per day, the opportunity to finally get commands into the lander would not be until 4 AM the next day). The lander receives the new commands during the 4 AM pass, and Earth finds out after the 4 PM pass if the new commands resolved the anomaly. At the earliest, assuming only one orbiter coverage, it would be on the order of 37 hours after an

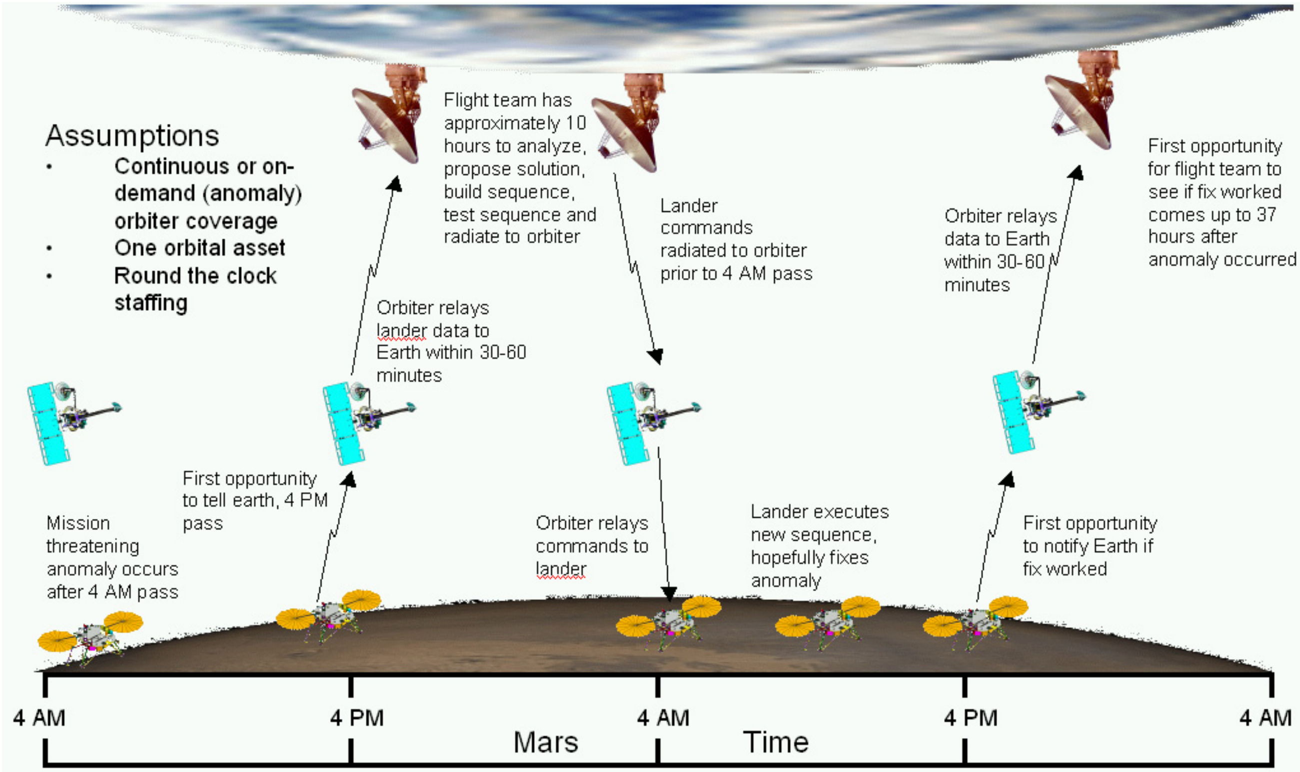

Figure 5. Worst case anomaly relay timeline showing long turn around for feedback and loss of landed operations time.

anomaly occurred before earth operators know if the solution was successful. This can be mission ending for certain types of power related problems. Even if the anomaly is not mission ending, an anomaly with a single instrument can mean a loss of at least two days of valuable landed science operations.

Energy and Power

As it turns out, power itself is not so critical as energy, that is power usage over time. The total energy collected in a sol is determined completely by the solar array size.
In the case of MPL and the '01 lander, a second order effect is the total load profile, which, because the lander uses a direct energy transfer system, affects the voltage across the arrays and thus the total power collected, so a
higher load means a lower voltage and less energy collected, but this affect is on the order of a few percent. One way to minimize this effect is to keep the total average load as constant as possible. The other limitation on total energy is the battery size, which limits how much of the collected energy can be used at night. Beyond those two limitations, current limits are the only other limiting factor. These considerations are taken into account in both the design phase, when resources are assigned, and in the operations planning phase.

On both MPL and the '01 lander, smaller (lower current rating) switches are grouped together downstream of larger (higher current rating) switches for several reasons. Quite often, the sum of the maximum current draw possible on each of the downstream switches is greater than capability of the upstream switch (for reasons of thermal design and simply quantity of copper). This can be managed however, by careful assignment of different components on the same upstream switch group. For example, payloads that are only operated (or only operated in their high power modes) on the surface can be grouped with flight system components that won't be used after landing, such as star trackers, IMU's, radars, or propulsion hardware. In the end, on the '01 lander, the switch assignments were such that after landing, almost all the payloads were able to use the max current rating of the upstream switch, albeit for short periods of time due to the energy limitations.

The energy limitations drive the design of the daily operations sequences. Data volume limitations also affect ops planning. For example, one sol might be considered a digging sol because of the high energy usage, or a sol might be an imaging sol because of the large data volume produced, and they might or might not be one and the same, depending on what else was planned. Some payloads have power spikes, such as from the operation of an oven for a few minutes or an hour, and other payloads needed to either be off or in lower power modes during these times. Figure 6 shows the resulting operations plan for the first 21 sols of the '01 lander mission.

6. Operations Approach

As landed payloads become more complex, so do the planning and the ultimate operations approach implemented.

For MPL, the operations approach was split into two phases. During cruise to Mars, most of the commanding was originated at the Lockheed Martin MSA in Denver. After landing, primary control was going to be from the payload center at UCLA, the home institution of the primary MPL payload, MVACS. As throughout the design, fabrication and test phase of the program, the same engineers who designed and built MPL helped operate it. The same people were also used for MCO, since both shared many subsystem commonalities, particularly in software and command and telemetry. This approach saved costs by allowing the same people to be used for two spacecraft, at least in areas where a whole person might not have otherwise been required. For the lander, once the spacecraft was on the surface, some of those people would no longer be required (the ACS and propulsion people, for example).

The solution for the landed mission was to provide on-sight support by a single systems engineer and general guidelines. Power and energy envelopes were generated and delivered to the payload operations center. If a payload profile did not exceed the parameters provided, the payload operators would know they were within the capabilities of the spacecraft. Ultimately, sequences still had to be run on the STL in Denver before they would be uplinked, but it reduced the overall turn-around time for planning. This would have been particularly helpful in the case of payload anomaly resolution.

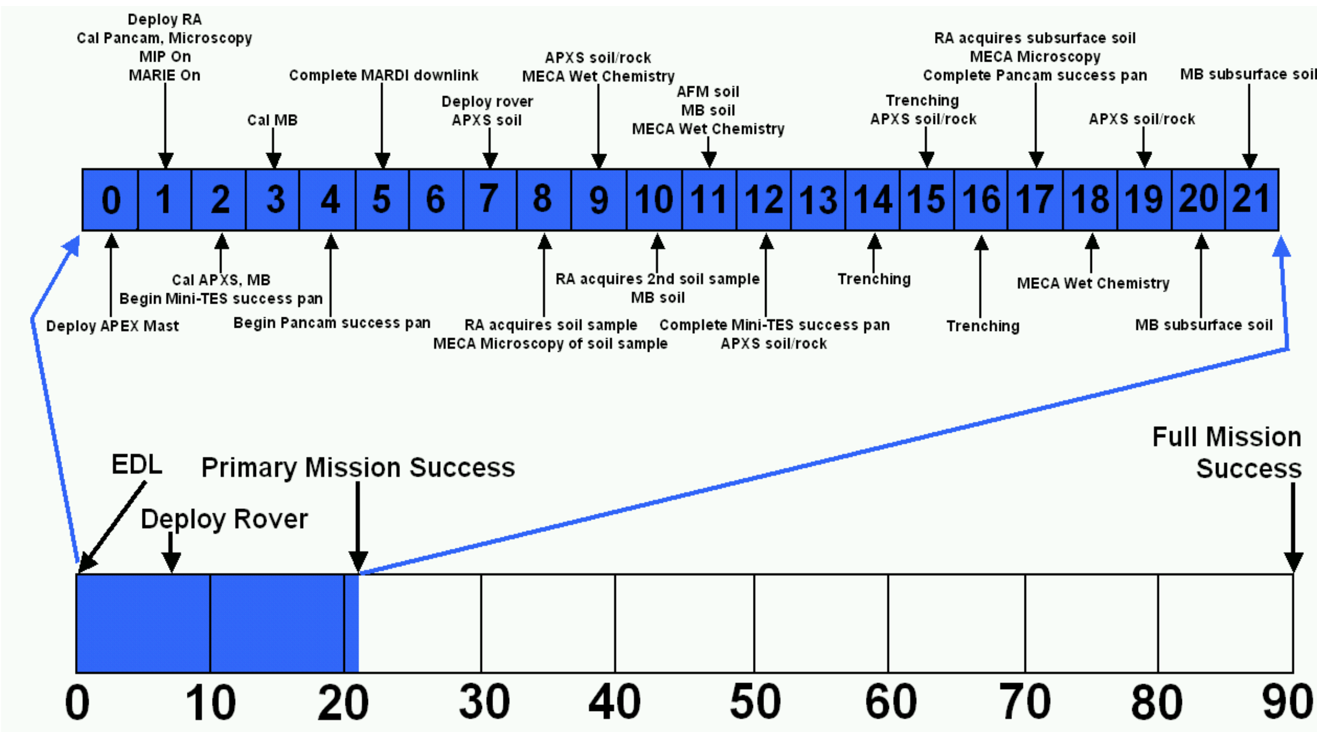

Figure 6. MSP 2001 lander 21 sol primary mission science campaign allocation developed with energy constraints.

For MSP'01, the landed operations center was to be at JPL, primarily because '01 was flying the Marie Curie rover and JPL had built an extensive experience base on Pathfinder with the Sojourner rover. The interaction between the payloads was even more extensive than on MPL. The robot arm was to be used to deploy the rover, provide soil samples to several instruments and had targets for the Athena instruments. PanCam was to be used to help rover navigation. It was to take images of experiments on several of the other payloads, as well as having it's own science objectives.

7. COST

A major benefit of the Mars Surveyor Program was the investment capitalization realized by producing a series of spacecraft. Each subsequent lander or orbiter benefited from the sum of experience and investment of the previous builds. While the '01 lander was a substantial upgrade from MPL, the basic design architecture was the same. This meant that lessons learned were more applicable than they might have otherwise been in a more independent design environment. Some efficiency of design and fabrication, often referred to as "the learning curve", were realized. While that learning curve also meant that some costs went up from what was originally bid, after the second time around, the costs associated with the process of designing, building and operating this style of spacecraft are even better understood. As a result, the MSP'01 project was able to realize a significant benefit from the money invested in the MSP'98 project. Because each project produced a lander and an orbiter using a high degree of synergy, it is not possible to completely separate the costs of either orbiter or lander. However, given the multiple builds to date of the Mars Surveyor spacecraft, there is a high degree of confidence in the ability to estimate the costs associated both with modifying and completing the '01 lander for flight and for building future spacecraft based on the same design. The loss of MPL and MCO prompted changes to the processes by which the Odyssey project was completed, however, this experience has been incorporated and is reflected in the current knowledge.

8. Conclusions

The evolutionary approach of developing landers begun by the Mars Surveyor program has resulted in designs for lightweight landers with payload fractions between 15%-20%. By leveraging the development effort of each of the previous vehicles, performance increases are possible that exceeded original expectations. The MPL payload capability was 24 kg, the ’01 lander capability is 66 kg, and with upgrades to the propulsion system, 120 kg of landed payload can be delivered to Mars. Other performance upgrades have been realized as well, as a result of the evolutionary approach, including increases in payload volume, power, energy storage and onboard resources. Operations approaches have also been developed to operate large, diverse payload suites.

9. Acknowledgments

This work was supported by the Mars Surveyor Program 2001, under contract to the Jet Propulsion Laboratory, contract number 960794, under contract to the National Aeronautics and Space Administration. The author would like recognize the MSP 2001 lander and MPL teams for their contributions to the spacecraft designs and operations.

10. References

[1] Casani, et al., “Report on the Loss of the Mars Polar Lander and Deep Space 2 Missions” *JPL Special Review Board*, JPL D-18709, March 22, 2000.

[2] E.A. Euler, S.D. Jolly, and H.H. Curtis, “The Failures of The Mars Climate Orbiter and Mars Polar Lander: A Perspective From the People Involved,” *24th Annual AAS Guidance and Control Conference*, Breckenridge, CO, January 31-February 4, 2001.

[3] J. Vellinga, “MPL Possible Premature Descent Engine Thrust Termination Process Investigation Report” *Lockheed Martin Astronautics Operations*, MSP-00-5001, March 10, 2000.

[4] Whetsel, et al., “Board Report for Return to Flight Review for Mars '01 Lander” *JPL Return to Flight Review Board*, IOM-313-CW-2000-01, March 1, 2000.

Richard Warwick *is a staff engineer at Lockheed Martin Astronautics. He has worked on JPL/NASA Mars missions since 1996. He received his M.S. and B.S. in aeronautical and astronautical engineering from the University of Washington in Seattle.*